



NSCL DAQ Compute/Storage

Tom Rockwell
FRIB IT Systems Group

MICHIGAN STATE

UNIVERSITY



U.S. DEPARTMENT OF
ENERGY

Office of
Science

This material is based upon work supported by the U.S. Department of Energy Office of Science under Cooperative Agreement DE-SC0000661, the State of Michigan and Michigan State University. Michigan State University designs and establishes FRIB as a DOE Office of Science National User Facility in support of the mission of the Office of Nuclear Physics.

Overview of Existing Infrastructure and Procedures

- NSF Data Policy
- DAQ Compute Resources
- DAQ Storage
- End-of-Experiment Procedures
- Recent Experience



Facility for Rare Isotope Beams
U.S. Department of Energy Office of Science
Michigan State University

NSF Data Policy

- NSCL policy / agreement between NSF and NSCL
- Key points related to IT Systems Group and DAQ systems:
 - “The **Spokesperson is responsible** for complying with NSF data policies, including long-term storage of the research data and records of the data analysis, and for responding to data access requests.”
 - “**NSCL will facilitate:** (1) the **recording of research data** during the running of the experiment, (2) the **transfer of the data to long-term storage media** by the Spokesperson or his/her designee, and will, as a courtesy, keep a duplicate of the recorded research data for a period of two years after completion of an experiment.”
- As a user facility, NSCL provides resources to assist with the proper recording and preserving of experiment data.



DAQ Compute Resources

- “DataU” - Experiment control room:
 - Several Experiment areas each with:
 - » 2 Windows desktop PCs
 - » 2 Linux Desktop PCs
 - » Additional resources as agreed
- “SPDAQ” Linux PCs running NSCLDaq software
 - Standard software and varying hardware configuration
 - » Debian Wheezy / Puppet configuration management
 - » DAQ subnet
 - » PCs: dual-core, 4 GB, SATA drive, 1GE
 - » Server PCs: dual-socket 16 GB, SATA drive, 1GE
- Available for duration of setup and experiment running period
 - Experiment account does not provide remote access
- Non-NSCL computers are not permitted on wired network



DAQ Storage Resources

- Linux/ZFS based NFS Servers
 - Data protection: parity, checksums, snapshots, etc.
 - Replication: enables online backups
- Systems:
 - “daqserv3”
 - » prioritized to running experiment
 - » ~ 20 TB (shared) */mnt/events/e12345*
 - » sized for \leq few TB per experiment
 - » sized for \leq few MB/s data recording rate
 - “daqserv4”
 - » backup system, hourly replication from daqserv3
 - NetApp
 - » */user/e12345* home directory



End-of-Experiment Procedure

- Once Spokesperson has signed form ...
- Experiment account is deactivated
- Events and home areas are transferred to “office” network
 - Available as /mnt/rawdata/e12345 (read-only)
 - » /mnt/events/e12345 -> /mnt/rawdata/e12345/events
 - » /user/e12345 -> /mnt/rawdata/e12345/home
 - retained online for 2 years
 - requires NSCL computer account for access
- LTO tape set written with contents of “rawdata”
 - Given to Spokesperson



Spring 2015, some fun – Data Rates

- Experiments recording high precision ADC output of waveforms
 - few 10MB/s writing
 - >TB per day
 - worked with experimenters to see that their needs met
- Successive experiments, nearly filling daqserv3 space
- Stressed operation procedures, but storage systems performed well
- Another doubling of data volume would be problematic with current systems



Spring 2015, some fun – semi-online Analysis

- Experiments analyzing “large” volume of data using office-side systems
 - “fishtank” faster shared Linux login interactive systems
 - “seaside” ~300 core batch system
- Experimenters (with NSCL user accounts)
 - “manually” moving data from DAQ to Office
 - and back
 - does it go to tape?
- System not designed with this activity in mind
 - interfering network flows
 - unexpected IO loads on storage
 - » potentially impacting data taking (was OK though)
- Will this activity increase, how to support ?
 - Potentially, IT replicates data to office side during experiment
 - Dedicated DAQ “offline” compute?

